

RADOS - Bug #24371

Ceph-osd crash when activate SPDK

06/01/2018 02:01 AM - Tone ZHANG

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Tone ZHANG		
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	luminous,mimic	Component(RADOS):	
Regression:	No	Pull request ID:	
Severity:	2 - major	Crash signature (v1):	
Reviewed:		Crash signature (v2):	

Description

Enable SPDK and configure bluestore as mentioned in <http://docs.ceph.com/docs/master/rados/configuration/bluestore-config-ref/>.

When launched the Ceph cluster with command "sudo MON=1 OSD=1 MDS=0 MGR=1 RGW=0 ../src/vstart.sh -n -x -l -b", met the ceph-osd crash as below:

```
2018-06-01 09:54:37.404 7fcb372b7200 -1 auth: unable to find a keyring on
/home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/keyring: (2) No such file or directory
2018-06-01 09:54:37.440 7fcb372b7200 -1 bluestore(/home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/block)
_read_bdev_label failed to open /home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/block: (2) No such file or directory
2018-06-01 09:54:37.440 7fcb372b7200 -1 bluestore(/home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/block)
_read_bdev_label failed to open /home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/block: (2) No such file or directory
2018-06-01 09:54:37.440 7fcb372b7200 -1 bluestore(/home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/block)
_read_bdev_label failed to open /home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0/block: (2) No such file or directory
2018-06-01 09:54:37.440 7fcb372b7200 -1 bluestore(/home/ubuntu/ceph-spdk/latest/20180514/ceph/build/dev/osd0) _read_fsid
unparsable uuid
Starting DDPK 17.11.0 initialization...
[ DDPK EAL parameters: nvme-device-manager -c 0x1 -m 2048 --file-prefix=spdk_pid29682 ]
EAL: Detected 20 lcore(s)
EAL: No free hugepages reported in hugepages-1048576kB
EAL: Probing VFIO support...
EAL: PCI device 0000:08:00.0 on NUMA socket 0
EAL: probe driver: 8086:953 spdk_nvme
```

- Caught signal (Segmentation fault)
in thread 7fcb372b7200 thread_name:ceph-osd
ceph version 13.1.0 (1f43eda5fd672d639637e539f6a5015ee215c8d2) mimic (dev)
1: (ceph::BackTrace::BackTrace(int)+0x45) [0x55630a3edf67]
2: (()+0x1cb3d50) [0x55630a69cd50]
3: (()+0x11390) [0x7fcb2ba9a390]
4: (std::__cxx11::List_base<aio_t, std::allocator<aio_t> >::M_clear()+0x2d) [0x55630a53145b]
5: (std::__cxx11::List_base<aio_t, std::allocator<aio_t> >::~List_base()+0x18) [0x55630a51f51c]
6: (std::__cxx11::list<aio_t, std::allocator<aio_t> >::~list()+0x18) [0x55630a515504]
7: (IOContext::~IOContext()+0x1e) [0x55630a519332]
8: (BlockDevice::reap_ioc()+0x1ae) [0x55630a625f44]
9: (SharedDriverQueueData::aio_handle(Task*, IOContext*)+0x12c7) [0x55630a68c4d5]
10: (NVMEDevice::aio_submit(IOContext*)+0x480) [0x55630a68fc40]
11: (NVMEDevice::read(unsigned long, unsigned long, ceph::buffer::list*, IOContext*, bool)+0x2ab) [0x55630a690983]
12: (BlueFS::_open_super()+0x195) [0x55630a62d239]
13: (BlueFS::mount()+0x10c) [0x55630a62bb2a]
14: (BlueStore::_open_db(bool, bool)+0x211d) [0x55630a4b4439]
15: (BlueStore::_fsck(bool, bool)+0x509) [0x55630a4c05c7]
16: (BlueStore::fsck(bool)+0x28) [0x55630a51a7c8]

```
17: (BlueStore::mkfs()+0x1c97) [0x55630a4bcd9]
18: (OSD::mkfs(CephContext*, ObjectStore*, std::__cxx11::basic_string<char, std::char_traits<char>,
std::allocator<char> > const&, uuid_d, int)+0xce) [0x556309d4298a]
19: (main()+0x1b84) [0x556309d14920]
20: (__libc_start_main()+0xf0) [0x7fcb2ac2e830]
21: (_start()+0x29) [0x556309d11ee9]
2018-06-01 09:54:41.224 7fcb372b7200 -1 Caught signal (Segmentation fault) *
in thread 7fcb372b7200 thread_name:ceph-osd
```

```
ceph version 13.1.0 (1f43eda5fd672d639637e539f6a5015ee215c8d2) mimic (dev)
1: (ceph::BackTrace::BackTrace(int)+0x45) [0x55630a3edf67]
2: (()+0x1cb3d50) [0x55630a69cd50]
3: (()+0x11390) [0x7fcb2ba9a390]
4: (std::__cxx11::_List_base< aio_t, std::allocator< aio_t > >::_M_clear()+0x2d) [0x55630a53145b]
5: (std::__cxx11::_List_base< aio_t, std::allocator< aio_t > >::~~_List_base()+0x18) [0x55630a51f51c]
6: (std::__cxx11::list< aio_t, std::allocator< aio_t > >::~~list()+0x18) [0x55630a515504]
7: (IOContext::~~IOContext()+0x1e) [0x55630a519332]
8: (BlockDevice::reap_ioc()+0x1ae) [0x55630a625f44]
9: (SharedDriverQueueData::_aio_handle(Task*, IOContext*)+0x12c7) [0x55630a68c4d5]
10: (NVMEDevice::aio_submit(IOContext*)+0x480) [0x55630a68fc40]
11: (NVMEDevice::read(unsigned long, unsigned long, ceph::buffer::list*, IOContext*, bool)+0x2ab) [0x55630a690983]
12: (BlueFS::_open_super()+0x195) [0x55630a62d239]
13: (BlueFS::mount()+0x10c) [0x55630a62bb2a]
14: (BlueStore::_open_db(bool, bool)+0x211d) [0x55630a4b4439]
15: (BlueStore::_fsck(bool, bool)+0x509) [0x55630a4c05c7]
16: (BlueStore::fsck(bool)+0x28) [0x55630a51a7c8]
17: (BlueStore::mkfs()+0x1c97) [0x55630a4bcd9]
18: (OSD::mkfs(CephContext*, ObjectStore*, std::__cxx11::basic_string< char, std::char_traits< char >,
std::allocator< char > > const&, uuid_d, int)+0xce) [0x556309d4298a]
19: (main()+0x1b84) [0x556309d14920]
20: (__libc_start_main()+0xf0) [0x7fcb2ac2e830]
21: (_start()+0x29) [0x556309d11ee9]
NOTE: a copy of the executable, or `objdump -rdS &lt;executable>` is needed to interpret this.
```

The crash is happened in Luminous and the latest version.

Related issues:

Copied to RADOS - Backport #24471: luminous: Ceph-osd crash when activate SPDK	Resolved
Copied to RADOS - Backport #24472: mimic: Ceph-osd crash when activate SPDK	Resolved

History

#1 - 06/01/2018 02:02 AM - Tone ZHANG

- Assignee set to Tone ZHANG

I'm working on the issue.

#2 - 06/01/2018 04:47 AM - Tone ZHANG

This is a bug in NVMEDevice, the bug fix has been committed.

Please have a review PR <https://github.com/ceph/ceph/pull/22356>

Thanks!

#3 - 06/04/2018 05:10 PM - Greg Farnum

- Project changed from Ceph to RADOS

- *Category deleted (OSD)*
- *Status changed from New to Fix Under Review*

#4 - 06/06/2018 12:33 AM - Kefu Chai

- *Status changed from Fix Under Review to Pending Backport*
- *Backport changed from luminous to luminous,mimic*

#5 - 06/09/2018 11:18 AM - Nathan Cutler

- *Copied to Backport #24471: luminous: Ceph-osd crash when activate SPDK added*

#6 - 06/09/2018 11:18 AM - Nathan Cutler

- *Copied to Backport #24472: mimic: Ceph-osd crash when activate SPDK added*

#7 - 10/03/2018 09:31 PM - Nathan Cutler

- *Status changed from Pending Backport to Resolved*